Isn't Hybrid ARQ Sufficient?

Michael Heindlmaier* and Emina Soljanin†
*Institute for Communications Engineering, Technische Universität München, Munich, Germany
†Bell Labs, Alcatel-Lucent, Murray Hill NJ 07974, USA
Email: michael.heindlmaier@tum.de, emina@alcatel-lucent.com

arXiv:1411.4061v1 [cs.IT] 14 Nov 2014

***Abstract*—In practical systems, reliable communication is often accomplished by coding at different network layers. We question the necessity of this approach and examine when it can be beneficial. Through conceptually simple probabilistic models (based on coin tossing), we argue that multicast scenarios and protocol restrictions may make concatenated multi-layer coding preferable to physical layer coding alone, which is mostly not the case in point-to-point communications.**

I. Introduction

A. Motivation

In packet-based data networks, large files are usually segmented into smaller blocks that are put into transport packets. Packet losses occur because of the physical channel and other limitations, such as processing power and buffer space. In current wireless systems, reliable communication is accomplished by coding at different network layers.

At the physical layer, a special transmission scheme, known as incremental redundancy Hybrid ARQ (IR-HARQ), which combines the conventional ARQ with error correction, has been in use since the appearance of 3G wireless technology (see, for example [1] for an overview). IR-HARQ schemes adapt their error code redundancy, based on the receiver's feedback, to varying channel conditions, and thus achieve better throughput performance than ordinary ARQ.

In broadcast/multicast applications from a single sender to many receivers, however, it is costly for the sender to collect and respond to individual receiver feedbacks, and thus HARQ schemes are disabled and packet losses are inevitable. With the rapid increase in multicast streaming applications, we see more and more proposals for packet level rateless erasure coding. A number of these schemes have already been standardized and are currently being implemented and deployed, e.g., Raptor codes for LTE eMBMS [2]. At the packet level, rateless codes enable efficient communications over multiple, unknown erasure channels, by asymptotically and simultaneously achieving the channel capacity at all erasure rates.

Although IR-HARQ, as a unicast technique, is disabled in multicast systems, physical layer coding at some chosen fixed rate remains. When this code fails to decode the noisy version of a data packet at the physical layer, the packet is declared erased, and data recovery is left to the packet level code. The rate of the physical layer code affects the successful-transmission time both positively and negatively. Increasing the rate, increases the number of channel uses, but lowers the packet erasure rate, which in turn means that fewer application layer rateless code packets will be sufficient for successful data transmission.

In practice, packet level coding schemes are not rateless. Standardized coding schemes start with a time-limited broadcast *delivery phase* by a content server using systematic Raptor codes with some fixed redundancy. This phase is often too short for all broadcast clients to collect a sufficient number of Raptor code symbols to be able to decode. Hence, once the delivery (broadcast) session expires, a unicast based file repair mechanism becomes available to the users who had experienced bad channel realizations and were not able to decode. These users enter the *repair phase*, during which the missing data is delivered by dedicated repair servers through unicasts with no packet level coding.

B. Related Work

Because of its practical relevance, this problem has been addressed in multiple ways, but only to a limited extent. Information theoretic analysis for single user scenarios suggest that channel codes be implemented entirely at the physical layer, where they can most efficiently combat fading [3], [4]. Furthermore, concatenated scheme involving coding at the physical as well as at a higher layer would be suboptimal.

This and related setups have also been investigated in [5]–[10]. Reference [5] studies the code rate tradeoff for both unicast and multicast setups where individual links are modeled as binary symmetric channels. Without packetization constraints, it turns out that in this case pure physical layer coding is optimal. In practice systems however, packetization is inevitable to some extent.

Reference [6] investigates the interplay between rate allocation and ARQ for point-to-point links, but do not consider coding at the packet level. Reference [7] investigates the tradeoff between physical layer and network layer rate allocation for networks in order to improve throughput, but does not study the possibility of HARQ. In an experimental setup, [8] studies the application of packet layer coding over a WiMAX point-to-point link: The best performance was observed with HARQ disabled and relying only on packet coding as reliability mechanism. The works [9], [10] compute the optimal physical layer code rate when packet level codes are operated in a rateless fashion, but do not study the impact of HARQ.

C. Our Approach

Our approach is to simplify the transmission model in order to better understand the interplay between the inner and the outer code for several single and multiuser scenarios that are motivated by current practice. The physical layer is modeled as a memoryless symbol erasure channel with fixed erasure probability ϵ_s which is given by the channel conditions. Successive symbol transmissions can thus be interpreted as flipping a biased coin, where the bias is given by system constraints and cannot be manipulated.

A packet contains k_s symbols and is successfully recovered if any k_s symbols are received, otherwise erased. The number of transmitted symbols, n_s, determines the coding rate of the physical layer. This rate can be adapted by means of HARQ, so n_s can vary from packet to packet in principle. For a fixed n_s successive packets are erased independently. Again, successive packet transmissions can thus be interpreted as flipping a biased coin where the bias can be adapted with the choice of n_s.

This simple coin tossing model is easy to understand but does not accurately represent the characteristics of the wireless channel. We still believe that the following essential feature is well approximated: In wireless systems with opportunistic link adaptation, the number of symbols needed to be transmitted for successful packet decoding is random because of imperfect channel knowledge. This is also captured by the coin tossing model.

The model does not capture the delay that is introduced by a practical implementation of HARQ schemes due to feedback delay or packet scheduling issues. These effects also have an impact on the delay performance [11]. The goal of this work is to show situations where, even without delay aspects, optimal coding operations become nontrivial.

This paper is organized as follows: In Sec. II, we present the system model under study. Different schemes for point-to-point scenarios are analyzed in Sec. III. We focus on the multiuser case in Sec. IV.

II. System Model and Problem Formulation

A. Data, Channel, and Code Models

We are transmitting a large data file of M chunks, each chunk has k_p packets, and each packet k_s channel symbols. Consequently, each chunk has $k_p \cdot k_s$ physical channel symbols. In transmission, the k_s data symbols of each packet may be protected by a channel code of length n_s at the physical layer, and the k_p data packets in each chunk may be protected by a packet level code of length n_p. This is visualized in Fig. 1.

The physical layer is modeled by a memoryless erasure channel with symbol erasure probability ϵ_s, assumed to be given. Coding is assumed to be such that a packet is successfully received if (any) k_s symbols are received, and a chunk is received if (any) k_p packets are received.

B. User, Redundancy, and Feedback Models

We will be considering *point-to-point* as well as *multicast* scenarios to u users. In a point-to-point scenario, either the inner or the outer code may be rateless [2], that is, coded symbols and packets may be sent until the packets and chunks are decoded. Note that if the inner (physical layer) code is rateless, the outer (packet level) code is not necessary.

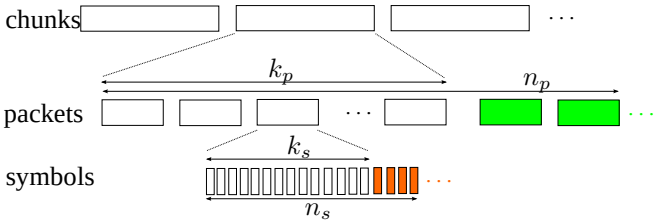

Fig. 1: System model for packetization

In multicast, the inner (physical channel) code may have a fixed rate, or the coded symbols may be transmitted until some ℓ users are able to decode. Note that $\ell \leq u$, and if $\ell = u$, then the outer (packet level) code is not needed.

We assume that each receiver can send a feedback signal at any moment on either level. The feedback is instantaneous, noiseless, and signifies only the completion of decoding. Under this assumption, the transmitter can send coded symbols (i.e., incremental redundancy) one symbol at the time, and does not have to send coded symbols that are not necessary for decoding.

C. Transmission Objectives

1) Expected File/Chunk Download Time: The goal of each user is to eventually download the file/chunk. Therefore, the most natural transmission objective is to minimize the expected time at which all user have successfully downloaded the file, which in general requires rateless transmission. In this paper, we will be mainly concerned with this objective, and our goal will be to find out at which level the rateless transmission should take place. Moreover, if the rateless coding is done only at the packet level, we are interested to find out which coding rate at the physical layer minimizes the overall multicast transmission time in this case.

2) Probability of Decoding Within Some Allocated Time: In practice, standardized coding schemes start with a time-limited broadcast delivery phase by a content server using systematic Raptor codes with some fixed redundancy. This phase is often too short for all broadcast clients to collect a sufficient number of Raptor code symbols to be able to decode. Therefore, another reasonable objective is to maximize probability of decoding under a delivery-time constraint, which in turn maximizes the expected number of users who are able to download the file in within this time and will not have to go through the file repair phase.

III. Point-to-Point Scenarios

Motivated by some common state-of-the-art schemes, we consider the following three point-to-point scenarios depending on the type of coding used at the (inner) physical layer:

1) Infinite Incremental Redundancy (IIR): This is an idealized rateless scheme where transmission at the physical layer of coded symbols is done until k_s symbols are received. Consequently, no packet gets erased, and thus no coding or retransmission at the packet level is needed.

2) Fixed Redundancy (FR): This scheme is motivated by the one in the LTE eMBMS protocol, where a fixed rate code (n_s, k_s) is used at the physical layer. Consequently, some packets will be erased, and coding or retransmission at the packet level is required.

3) Finite Incremental Redundancy (FIR): This scheme most closely resembles the current IR-HARQ implementations. Here again, a fixed rate code (n_s, k_s) is used at the physical layer, and transmission of coded symbols is done until either k_s symbols are received or n_s coded symbols are used up (whichever happens first). Consequently, some packets will be erased, and coding or retransmission at the packet level is required.

Recall that, since our objective is to ultimately decode the chunk, either the inner or the outer code must be rateless, that is, must allow unlimited number of transmissions until decoding can be performed. Let T_i^s, be the number of channel symbols used for the i-th packet transmission, and T^p the number of (coded) packets that have to be sent to decode the chunk. Then decoding of the chunk requires

$$\sum_{i=1}^{T^p} T_i^s \tag{1}$$

channel transmissions. Note that, depending on the scenario, either T_i^s or T^p or both can be random variables.

We next describe the three point-to-point scenarios, derive the expected number of channel transmissions necessary for decoding, and provide some quantitative examples. We conclude the section by showing that using rateless coding at the inner rather than at the outer layer results in fewer number of transmissions on average.

A. Infinite Incremental Redundancy (IIR) Schme

In the IIR scheme, the code is rateless at the physical layer, meaning that coded symbols are transmitted over the channel with the erasure probability ϵ_s until k_s of them are received, at which point the receiver can decode the data. Therefore, the number of transmissions is a negative binomial random variable with parameters $(k_s, 1-\epsilon_s)$[1], which we will denote by $NB(k_s, 1-\epsilon_s)$ The number of channel transmissions required to successfully decode $T^p = k_p$ packets is the sum of these random variables, and thus itself a negative binomial random variable with parameters $(k_p \cdot k_s, 1-\epsilon_s)$. We denote this time by T^{IIR}. Its expected value is given by

$$\mathbb{E}\big[T^{\text{IIR}}\big] = \frac{k_p \cdot k_s}{1-\epsilon_s} \tag{2}$$

[1] $1-\epsilon_s$ is the probability of success in the associated Bernoulli trials.

B. Fixed Redundancy (FR) Scheme

In the FR scheme, a fixed-rate code (n_s, k_s) is used at the physical layer, and thus it always takes $T_i^s = n_s$ channel transmissions per data packet, at which point the packet is either successfully decoded (when k_s or more channel transmissions are successful) or erased with probability ϵ_p, given by

$$\epsilon_p = \sum_{j=k_s+1}^{n_s} \binom{n_s}{j} \epsilon_s^j (1-\epsilon_s)^{(n_s-j)}. \tag{3}$$

The FR scheme is rateless at the packet level. Therefore, the number of packet transmissions over the channel with the erasure probability ϵ_p until k_p transmissions are successful (at which point the receiver can decode the data) is a negative binomial random variable with parameters $(k_p, 1-\epsilon_p)$. Since each packet transmission takes exactly n_s channel transmissions, the number of channel transmissions T^{FR} required to decoded the file is not a negative binomial random variable, but its expected value is given by

$$\mathbb{E}\big[T^{\text{FR}}\big] = \frac{k_p \cdot n_s}{1-\epsilon_p}. \tag{4}$$

Note that it is not immediately clear how (4) compares to (2) since $n_s \geq k_s$ but $\epsilon_s \geq \epsilon_p$. However, we show below that $\mathbb{E}\big[T^{\text{FR}}\big] \geq \mathbb{E}\big[T^{\text{IIR}}\big]$, that is, using rateless coding at the inner rather than at the outer layer results in fewer number of transmissions on average.

C. IIR vs. FR

Theorem 1. $\mathbb{E}\big[T^{FR}\big] \geq \mathbb{E}\big[T^{IIR}\big]$

Proof: Consider the non-negative random variable S corresponding to the number of successfully received symbols after n_s transmission. Note that S is binomial with parameters n_s and $1-\epsilon_s$, and therefore its expectation is $n_s \cdot (1-\epsilon_s)$. By the definition of ϵ_p and the Markov's inequality (see e.g., [12, p. 116]), we have

$$1-\epsilon_p = P(S \geq k_s) \leq \frac{n_s \cdot (1-\epsilon_s)}{k_s}, \tag{5}$$

which gives

$$\frac{k_s}{1-\epsilon_s} \leq \frac{n_s}{1-\epsilon_p},$$

and the claim follows from (2) and (4). ■

Therefore rateless coding at the inner rather than at the outer layer results in fewer channel uses on average for chunk download. However, when ϵ_s is known, then n_s can be optimized to minimize $n_s/(1-\epsilon_p)$, then $\mathbb{E}\big[T^{\text{IIR}}\big]$ and $\mathbb{E}\big[T^{\text{FR}}\big]$ do not differ much, as shown by an example in Fig. 2.

Note instantaneous and perfect feedback should be possible after each symbol for IIR, and therefore, when the channel is known and stays constant, FR beats IIR in practice. However, in current practical wireless scenarios, the channel is unpredictable, and a HARQ scheme in use is similar to the one we consider in the next section.

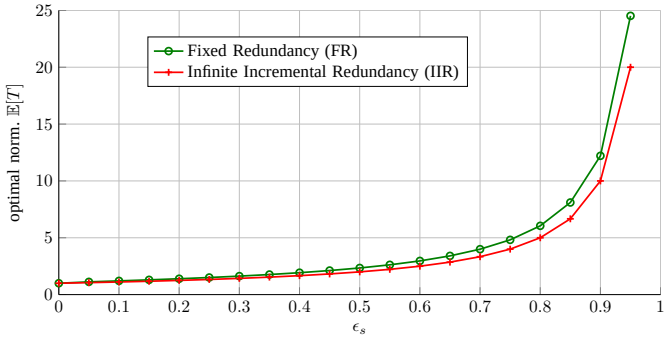

Fig. 2: Normalized delay for FR and IIR for $k_s = 100$. n_s was chosen to minimize $\mathbb{E}[T^{\text{FR}}]$.

D. Finite Incremental Redundancy (FIR) Scheme

In the FIR scheme, both T^p or T_i^s are random variables in contrast to the previous schemes: At the physical layer at most n_s symbols are transmitted for each transmission, so $k_s \leq T_i^s \leq n_s$. The associated erasure probability ϵ_p is as in (3). Let X_i denote the number of symbols that would be needed to receive the i-th transmitted packet correctly. X_i is distributed according to a negative binomial distribution with parameters k_s and $1-\epsilon_s$. The random variable T_i^s is now given by $T_i^s = \min\{X_i, n_s\}$. As the sequence $(T_i^s)_{i=1}^{\infty}$ is i.i.d., one can use Wald's generalized equation[2] [13, Theorem 5.5.2] to show that

$$\mathbb{E}\big[T^{\text{FIR}}\big] = \mathbb{E}\big[T^p\big]\mathbb{E}\big[T_i^s\big]. \tag{6}$$

$\mathbb{E}\big[T^p\big]$ is the expected value of a negative binomial random variable with parameters k_p and $1-\epsilon_p$. The second expected value is given by

$$\begin{aligned}\mathbb{E}[T_i^s] = &\sum_{n=k_s}^{n_s} n\binom{n-1}{k_s-1}(1-\epsilon_s)^{k_s}\epsilon_s^{n-k_s} \\ &+ n_s \sum_{n=n_s+1}^{\infty}\binom{n-1}{k_s-1}(1-\epsilon_s)^{k_s}\epsilon_s^{n-k_s}\end{aligned} \tag{7}$$

and cannot be solved in closed form. Hence,

$$\mathbb{E}\big[T^{\text{FIR}}\big] = \frac{k_p}{1-\epsilon_p}\mathbb{E}\big[T_i^s\big]. \tag{8}$$

As for FR, the performance of the FIR scheme depends on the choice of n_s. Choosing n_s too small will result in a high packet erasure probability ϵ_p. Letting $n_s \to \infty$ the scheme approaches the performance of the IIR scheme. This is also shown in Fig. 3, which visualizes the dependency of the expected delay on the choice of n_s.

IV. Multicast Scenarios

We are now concerned with multicast to u users, each independently experiencing a channel with erasure probability ϵ_s. As in the point-to-point case, we distinguish between transmission scenarios depending on the type of coding used at the (inner) physical layer. Again, since the objective is ultimate decoding of the file (this time by all users), rateless coding is necessary at either physical or packet level, but now coded symbols or packets have to be transmitted until all users are able to decode. We first analyze and then compare the IIR and FR schemes.

[2]The general version of Wald's equation should be used because T^p is a stopping time that cannot directly be defined as a function of $(T_i^s)_{i=1}^{\infty}$.

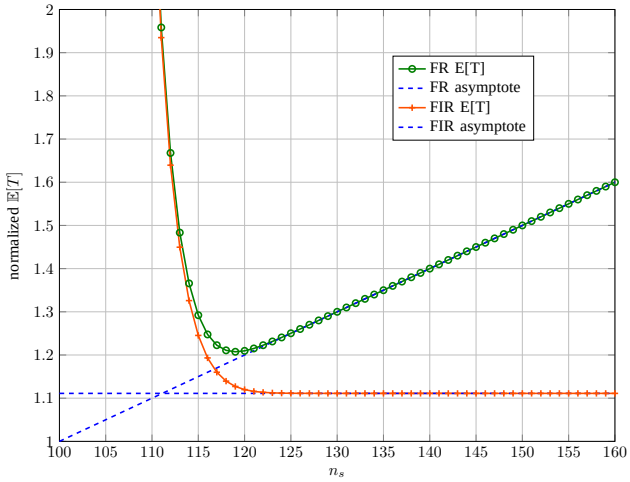

Fig. 3: Normalized expected delay for FR and FIR as a function of n_s for $\epsilon_s = 0.1$ and $k_s = 100$.

A. IIR with u Users

Let $T_i^s(j)$ be the number of coded symbol-transmissions until user j, $j \in \{1, \ldots, u\}$ is able to decode packet i, $i \in \{1, \ldots, k_p\}$. Note that $T_i^s(j)$ are i.i.d. $NB(k_s, 1-\epsilon_s)$. Then the number of coded symbol-transmissions until packet i is decoded by all users, is given by

$$T_i^s = \max_{j\in\{1,\ldots,u\}} T_i^s(j),$$

and is therefore distributed as the maximum order statistic of u of i.i.d. $NB(k_s, 1-\epsilon_s)$.

In the IIR scheme, multicast of packet i starts only when all users decode packet $i-1$, and a chunk is decoded when all packets are decoded by all users. Therefore, the number of symbol-transmissions until a chunk is decoded is given by

$$T^{\text{IIR}} = \sum_{i=1}^{k_p} T_i^s = \sum_{i=1}^{k_p} \max_{j\in\{1,\ldots,u\}} T_i^s(j)$$

and its average by

$$\mathbb{E}\big[T^{\text{IIR}}\big] = k_p \cdot \mathbb{E}\big[T_1^s\big].$$

B. FR with u Users

In the FR scheme, each packet takes exactly n_s symbol transmissions. At that point some users will not be able to decode the packet, and that will happen with probability ϵ_p given by (3). Let $T^p(j)$ be the number of coded packet-transmissions until user j, $j \in \{1, \ldots, u\}$ is able to decode the chunk. Note that $T^p(j)$ are i.i.d. $NB(k_p, 1-\epsilon_p)$. Therefore,

the number of packet-transmissions until a chunk is decoded is given by

$$T^p = \max_{j\in\{1,\ldots,u\}} T^p(j)$$

and the average number of symbol-transmissions until a chunk is decoded is given by

$$\mathbb{E}\big[T^{\text{FR}}\big] = n_s \cdot \mathbb{E}\big[T_p\big]$$

C. IIR vs. FR with u Users

We have seen that in the point-to-point case, coding rateless at the symbol level results in fewer transmissions on average. As we will see, that is not always the case for the multiple users. Earlier in this section, we have seen the following:

- The IIR scheme chunk download time (measured by the number of symbol-transmissions) is equal to the "maximum order statistic of u of $NB(k_s, 1-\epsilon_s)$" $\times$ k_p.
- The FR scheme chunk download time (measured by the number of symbol-transmissions) is equal to the "maximum order statistic of u of $NB(k_p, 1-\epsilon_p)$" $\times$ n_s.

The expected value of the maximum order statistic of negative binomial random variables with parameters $(k, 1-\epsilon)$ has been investigated in [14]. One result of this work is the following approximation:

$$\log_{1/\epsilon} u + (k-1)\log_{1/\epsilon}\big[\log_{1/\epsilon} u\big]$$

We found in our numerical simulations that this approximation is not precise enough in many cases. Instead, we found that the expectation of the maximum order statistic of u of $NB(k, 1-\epsilon)$ behaves as follows (see Appendix for an argument):

- A good numerical approximation when k is small is

  $$\log_{1/\epsilon} u + (k-1)\big[\log_{1/\epsilon}\log_{1/\epsilon} u + \log_{1/\epsilon}(1-\epsilon)/\epsilon\big]$$

  which also gives good qualitative description for larger k that can be used to get an insight into the behaviour (if not the precise value) of the chunk download time for each of the schemes.
- $k/(1-\epsilon)$ is a good approximation when k is very large or when k is large compared to u, and also when ϵ is very small (but for different probabilistic reasons than when k is large, see Appendix).

Our approximation indicates that the expected value of u of $NB(k, 1-\epsilon)$ will

1) increase with u,
2) have a lower rate of increase with u as k gets larger,
3) have a highr rate of increase with u as ϵ gets larger.

Figures 4 and 5 obtained by simulation clearly illustrate this behaviour.

As can be seen from the figures, for a single user or a small number of users, the IIR scheme has a lower expected download time. This encourages the use of Hybrid-ARQ only. However, for a larger group of users, it is better to admit a certain packet erasure probability and rely on packet level coding for reliable transmission.

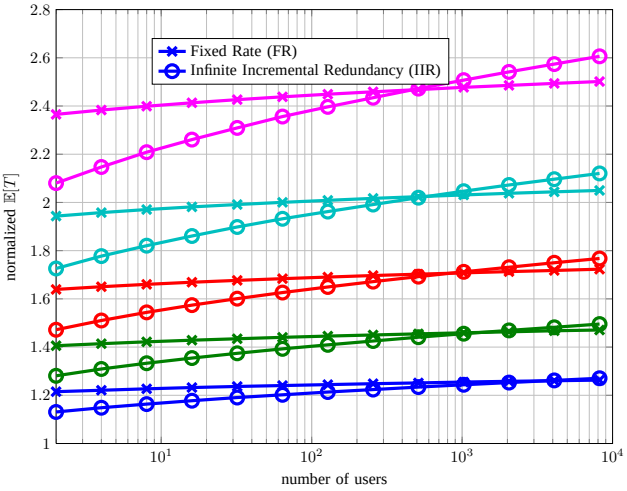

Fig. 4: Normalized expected delay for FR and IIR as a function of the number of users for $k_p = 100$ and $k_s = 100$, $\epsilon_s \in \{0.1, 0.2, 0.3.0.4, 0.5\}$.

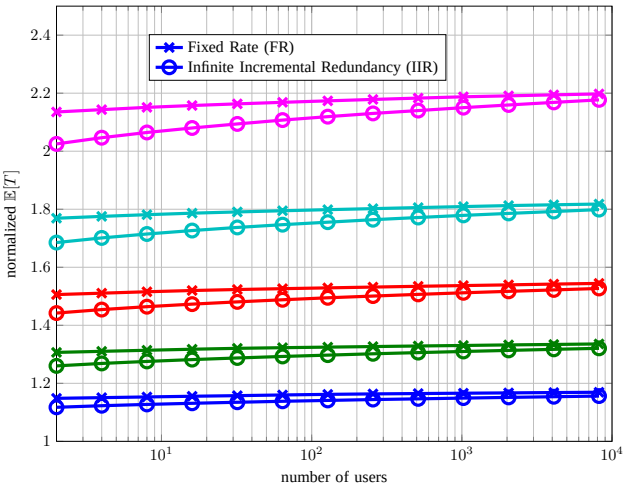

Fig. 5: Normalized expected delay for FR and IIR as a function of the number of users for $k_p = 100$ and $k_s = 1000$, $\epsilon_s \in \{0.1, 0.2, 0.3.0.4, 0.5\}$.

This result might be counter-intuitive, but can be explained in the following way: Both the expected download time of IIR and FR involve the expected value of the maximum of negative binomial random variables. Only the parameters are different. It is thus not immediately clear why the performance curves show the observed behavior. One possible explanation is as follows: For a large number of users, the time needed to satisfy *all* u users in the IIR scheme depends on the value of ϵ_s. Similarly, in the FR scheme, the number of packets needed to satisfy all u users depends on ϵ_p. The value of ϵ_p can be influenced through the optimal choice of n_s whereas ϵ_s cannot be changed. We have thus the possibility to adapt the coin flipping bias associated with the packet transmission

- and this leads to a smaller expected download time for FR as the number of users large number of users gets large.

V. Conclusion and Further Work

We investigated the relationship of Hybrid-ARQ and rateless packet level coding schemes. Using a coin tossing model we observe the well-accepted fact that Hybrid-ARQ alone is advantageous in point-to-point scenarios. However, for multicast scenarios with many users, we observe that a packet level coding can outperform schemes that rely on Hybrid-ARQ only. In this context, we reviewed approximations for order statistics of negative binomial random variables and proposed improved approximations.

For future work, we target to extend these results to more general cases and add refined model assumptions.

Appendix

We here study the expected value of the maximum order statistic of u of $NB(k, 1-\epsilon)$. Consider a game with u players, each independently tossing a biased coin with the probability of head equal to $(1-\epsilon)$ until he sees k heads. We separately analyze three cases:

1) Small k, Large u: The probability that a players did not see any heads after t tosses is equal to ϵ^t. Therefore, the expected number of such players is $u \cdot \epsilon^t$, which is smaller than 1 when $t > t_1$, where

$$t_1 = \log_{1/\epsilon} u.$$

The probability that a players saw only one head after t tosses is equal to $t \cdot \epsilon^{(t-1)}(1-\epsilon)$. Therefore, the expected number of such players is $u \cdot t \cdot \epsilon^{(t-1)}(1-\epsilon)$, which becomes equal to $t_1 \cdot (1-\epsilon)/\epsilon$ when $t = t_1$. By the same argument as before, these players need another $\Delta t = \log_{1/\epsilon} t_1 \cdot (1-\epsilon)/\epsilon$ tosses on average until each of these users has seen at least one more head on average. Therefore, after $t_1 + \log_{1/\epsilon} t_1 \cdot (1-\epsilon)$ tosses the average, the number of users who have seen fever than 2 heads is smaller than 1.

Let t_k denote the moment at which the expected number of players who have seen k heads is smaller than 1. By the above reasoning, we have that

$$t_2 = t_1 + \Delta t = \log_{1/\epsilon} u + \log_{1/\epsilon} \log_{1/\epsilon} u + \log_{1/\epsilon}(1-\epsilon)/\epsilon$$

and for small k, we have

$$t_k \approx \log_{1/\epsilon} u + (k-1)\big[\log_{1/\epsilon} \log_{1/\epsilon} u + \log_{1/\epsilon}(1-\epsilon)/\epsilon\big].$$

2) Large k: For large k, the number of tosses that the players have to make is also large, and by the law of large numbers, each player will take approximately $k/(1-\epsilon)$ to see k heads.

3) Small ϵ: When ϵ is small, then $1-\epsilon$ is large, and thus almost all users see a head each time they toss their coins. Note that, as predicted by 2) and 3) above, the dependence of the download time on the number of users is small when k is large or ϵ is small.

Figure 6 illustrates the tightness of our bound and the bound of [14].

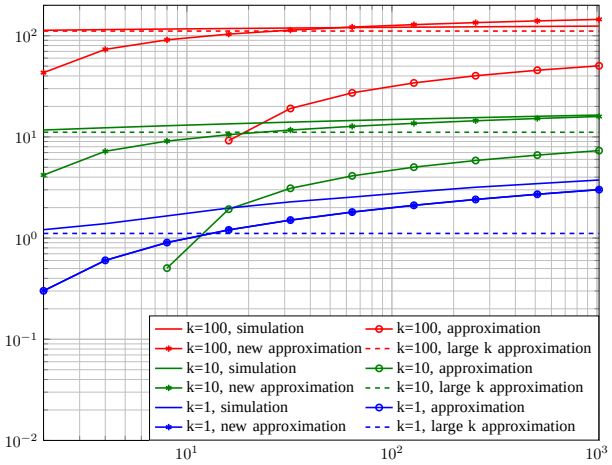

Fig. 6: Order statistics approximations and simulation results for $\epsilon = 0.1$. Both approximations are shown for different values of k. Dashed lines show the approximation that holds if $k \gg u$. For $k = 1$, both approximations match.

References

[1] C. Lott, O. Milenkovic, and E. Soljanin, "Hybrid ARQ: Theory, state of the art and future directions," in *IEEE Inf. Theory Workshop (ITW)*, Bergen, Norway, Jul. 2007.

[2] M. A. Shokrollahi and M. Luby, "Raptor codes," *Foundations and Trends in Communications and Information Theory*, vol. 6, no. 3-4, 2009.

[3] C. R. Berger, S. Zhou, Y. Wen, P. Willett, and K. Pattipati, "Optimizing joint erasure-and error-correction coding for wireless packet transmissions," *IEEE Trans. Wireless Comm.*, vol. 7, no. 11, pp. 4586–4595, 2008.

[4] T. A. Courtade and R. D. Wesel, "Optimal allocation of redundancy between packet-level erasure coding and physical-layer channel coding in fading channels," *IEEE Trans. Comm.*, vol. 59, no. 8, pp. 2101–2109, 2011.

[5] C. Koller, M. Haenggi, J. Kliewer, and D. Costello, "On the optimal block length for joint channel and network coding," in *IEEE Inf. Theory Workshop (ITW)*. IEEE, 2011, pp. 528–532.

[6] P. Wu and N. Jindal, "Coding versus ARQ in fading channels: How reliable should the PHY be?" *IEEE Trans. Comm.*, vol. 59, no. 12, pp. 3363–3374, December 2011.

[7] T. Cui, T. Ho, and J. Kliewer, "Achievable rate and optimal physical layer rate allocation in interference-free wireless networks," in *IEEE Int. Symp. Inf. Theory*. IEEE, 2009, pp. 2674–2678.

[8] S. Teerapittayanon, K. Fouli, M. Médard, M.-J. Montpetit, X. Shi, I. Seskar, and A. Gosain, "Network coding as a WiMAX link reliability mechanism," in *Multiple Access Communications*. Springer, 2012, pp. 1–12.

[9] M. Vehkapera and M. Médard, "A throughput-delay trade-off in packetized systems with erasures," in *IEEE Int. Symp. Inf. Theory (ISIT)*. IEEE, 2005, pp. 1858–1862.

[10] M. Xiao, M. Médard, and T. Aulin, "Cross-layer design of rateless random network codes for delay optimization," *IEEE Trans. Comm.*, vol. 59, no. 12, pp. 3311–3322, 2011.

[11] A. Larmo, M. Lindstrom, M. Meyer, G. Pelletier, J. Torsner, and H. Wiemann, "The LTE link-layer design," *IEEE Communications Magazine*, vol. 47, no. 4, pp. 52–59, 2009.

[12] P. Flajolet and R. Sedgewick, *Analytic Combinatorics*, 1st ed. New York, NY, USA: Cambridge University Press, 2009.

[13] R. G. Gallager, *Course notes of Discrete stochastic processes*. Kluwer Academic Publishers Boston, 2013, vol. 101.

[14] P. J. Grabner and H. Prodinger, "Maximum statistics of n random variables distributed by the negative binomial distribution," *Combinatorics, Probability and Computing*, vol. 6, no. 02, pp. 179–183, 1997.